Notice of Allowability	Application No.	Applicant(s)	
	10/017,337	CAO ET AL.	
	Examiner	Art Unit	
	Kandasamy Thangavelu	2123	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☒ This communication is responsive to *June 23, 2005*.

2. ☒ The allowed claim(s) is/are *11,13,14,16-22,27-30,34,35 and 37-63*.

3. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
   a) ☐ All b) ☐ Some* c) ☐ None of the:
      1. ☐ Certified copies of the priority documents have been received.
      2. ☐ Certified copies of the priority documents have been received in Application No. _____ .
      3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

   * Certified copies not received: _____.

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

4. ☐ A SUBSTITUTE OATH OR DECLARATION must be submitted. Note the attached EXAMINER'S AMENDMENT or NOTICE OF INFORMAL PATENT APPLICATION (PTO-152) which gives reason(s) why the oath or declaration is deficient.

5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
   (a) ☐ including changes required by the Notice of Draftsperson's Patent Drawing Review (PTO-948) attached
      1) ☐ hereto or 2) ☐ to Paper No./Mail Date _____.
   (b) ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____.

   Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).

6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

1. ☒ Notice of References Cited (PTO-892)
2. ☐ Notice of Draftperson's Patent Drawing Review (PTO-948)
3. ☒ Information Disclosure Statements (PTO-1449 or PTO/SB/08), Paper No./Mail Date 23 June 2005
4. ☐ Examiner's Comment Regarding Requirement for Deposit of Biological Material
5. ☐ Notice of Informal Patent Application (PTO-152)
6. ☐ Interview Summary (PTO-413), Paper No./Mail Date _____ .
7. ☒ Examiner's Amendment/Comment
8. ☒ Examiner's Statement of Reasons for Allowance
9. ☒ Other *Clean copy of allowed claims*.

DETAILED ACTION

Introduction

1. This communication is in response to the Applicants' communication dated June 23, 2005. Claims 1-10 were canceled. Claims 11,14, 15, 24, 25 and 26 were amended. Claims 30-39 were added. Claims 11-39 of the application are pending.

Information Disclosure Statement

2. Acknowledgment is made of the information disclosure statements filed on June 23, 2005 together with a list patents and papers. The patents and papers have been considered.

Examiner's Amendment

3. Authorization for this examiner's amendment was given in a telephone conversation by Mr. Omkar Suryadevara on September 15, 2005.

An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to the applicants, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

3. In the claims:

Replace Claim 11 with:

11. A computer-readable storage medium having stored thereon computer instructions that, when executed by a computer, cause the computer to simulate effect of a recovery time required for the database system to recover from a database failure on runtime performance of the database system, the instructions comprising instructions to:

provide in the computer, in addition to a normal checkpoint queue used in the database system for normal operation, at least one simulated checkpoint queue;

wherein the normal checkpoint queue comprises a plurality of buffers;

wherein the simulated checkpoint queue is an ordered list of one or more elements, each element in the simulated checkpoint queue representing a respective buffer that is or was in the normal checkpoint queue, the ordered list having a head and a tail;

wherein the simulated checkpoint queue is associated with a setting for simulated mean time to recover (MTTR) whose effect on runtime performance of the database system is being simulated in the computer;

in response to detecting a change to a buffer in the normal checkpoint queue due to actual database transactions occurring within the database system under normal operating conditions, check if the buffer is represented in the simulated checkpoint queue, and if the buffer is not represented in the simulated checkpoint queue, link an element that represents the buffer to the tail of the simulated checkpoint queue;

provide a simulated write counter, the simulated write counter being associated with the simulated MTTR setting;

wherein the simulated write counter provides a count of a number of times any element is removed from the simulated checkpoint queue in response to a simulated write out of the respective buffer from volatile memory to nonvolatile memory;

determine if linking the element to the tail of the simulated checkpoint queue causes the simulated checkpoint queue to exceed a predetermined length; and

in response to determining that the simulated checkpoint queue exceeds the predetermined length, remove an element from the head of the simulated checkpoint queue and increment the simulated write counter.

In Claim 12:

Delete Claim 12.

Replace Claim 14 with:

14. The computer-readable storage medium of Claim 11 further storing computer instructions that, when executed by a computer, cause the computer to, write out of the any buffer from volatile memory and storing in nonvolatile memory using an incremental checkpoint operation in the normal operation of the database;

in response to detecting a write out of any buffer by the incremental checkpoint operation, check if that buffer is represented in the simulated checkpoint queue, and if that buffer is represented in the simulated checkpoint queue, remove the element

representing that buffer from the simulated checkpoint queue and increment the simulated write counter.

In Claim 15:

Delete Claim 15.

In Claim 16, Line 1, "The computer-readable storage medium of Claim 10"

Has been changed to

-- The computer-readable storage medium of Claim 11--.

Replace Claim 21 with:

21. A system for simulating effect of a recovery time required for the database system to recover from a database failure on runtime performance of the database system comprising:

a memory for storing program instructions and data;

one or more processors coupled to the memory;

a simulated MTTR setting maintained in the memory;

a normal checkpoint queue used in the database system for normal operation, wherein the normal checkpoint queue comprises a plurality of buffers;

at least one simulated checkpoint queue;

wherein the simulated checkpoint queue is an ordered list of one or more elements, each element in the simulated checkpoint queue representing a respective

buffer that is or was in the normal checkpoint queue, the ordered list having a head and a tail;

wherein the simulated checkpoint queue is associated with a setting for simulated mean time to recover (MTTR) whose effect on runtime performance of the database system is being simulated in the computer; and

a simulated write counter maintained in the memory, the simulated write counter being associated with the simulated MTTR setting;

wherein the simulated write counter provides a count of the number of times an element is removed from the simulated checkpoint queue, in response to a simulated write out of a buffer from volatile memory and storing in nonvolatile memory;

wherein in response to detecting a change to a buffer in the normal checkpoint queue due to actual database transactions occurring within the database system under normal operating conditions, the instructions in memory check if the buffer is represented in the simulated checkpoint queue, and if the buffer is not represented in the simulated checkpoint queue, the instructions link an element that represents the buffer to the tail of the simulated checkpoint queue;

wherein the instructions in memory determine if linking the element to the simulated checkpoint queue causes the simulated checkpoint queue to exceed a predetermined length; and

in response to determining that the simulated checkpoint queue exceeds the predetermined length, the instructions in memory remove an element from the head of the simulated checkpoint queue and increment the simulated write counter.

Replace Claim 22 with:

22. The system of Claim 21, wherein instructions in memory write out of the any buffer from volatile memory and store in nonvolatile memory using an incremental checkpoint operation in the normal operation of the database;

in response to detecting a write out of any buffer by the incremental checkpoint operation, instructions in memory check if that buffer is represented in the simulated checkpoint queue, and if that buffer is represented in the simulated checkpoint queue, remove the element representing that buffer from the simulated checkpoint queue and increment the simulated write counter.

In Claims 23-26:

Delete claims 23-26.

Replace Claim 30 with:

30. A computer implemented method that comprises a database system, to simulate effect of a recovery time required for the database system to recover from a database failure on runtime performance of the database system, the method comprising:

providing in the computer, in addition to a normal checkpoint queue used in the database system for normal operation, at least one simulated checkpoint queue;

wherein the normal checkpoint queue comprises a plurality of buffers;

wherein the simulated checkpoint queue is an ordered list of one or more elements, each element in the simulated checkpoint queue representing a respective

buffer that is or was in the normal checkpoint queue, the ordered list having a head and a tail;

wherein the simulated checkpoint queue is associated with a setting for recovery time whose effect on runtime performance of the database system is being simulated in the computer;

in response to detecting a change to a buffer in the normal checkpoint queue due to actual database transactions occurring within the database system under normal operating conditions, checking if the buffer is represented in the simulated checkpoint queue, and if the buffer is not represented in the simulated checkpoint queue, linking an element that represents the buffer to the tail of the simulated checkpoint queue;

providing a simulated write counter, the simulated write counter being associated with the setting for recovery time;

wherein the simulated write counter provides a count of a number of times any element is removed from the simulated checkpoint queue in response to a simulated write out of the respective buffer from volatile memory to nonvolatile memory;

determining if linking the element to the simulated checkpoint queue causes the simulated checkpoint queue to exceed a predetermined length; and

in response to determining that the simulated checkpoint queue exceeds the predetermined length, removing an element from the head of the simulated checkpoint queue and incrementing the simulated write counter.

In Claims 31-33:

Delete claims 31- 33.

Replace Claim 35 with:

35. The method of Claim 30 further comprising:

writing out of the any buffer from volatile memory and storing in nonvolatile memory using an incremental checkpoint operation in the normal operation of the database;

in response to detecting a write out of any buffer by the incremental checkpoint operation, checking if that buffer is represented in the simulated checkpoint queue, and if that buffer is represented in the simulated checkpoint queue, removing the element representing that buffer from the simulated checkpoint queue and incrementing the simulated write counter.

In Claim 36:

Delete Claim 36.

Add claims 40 to 63 as follows:

40. A computer-implemented method, to simulate effect on runtime performance of a database system, of a plurality of settings to control mean time to recover (MTTR) from a database failure, the method comprising:

using a current setting to limit MTTR of said database system for normal operation; and

simulating performance of said database system for a plurality of additional settings of MTTR, using data from said normal operation.

41. The method of Claim 40, wherein said data comprises actual database operating data from said database system.

42. The method of Claim 40 wherein:

said simulating is performed concurrently with said performing of normal operation.

43. The method of Claim 42 wherein:

results of said simulating are available in real-time.

44. The method of Claim 42 wherein:

results of said simulating are available in quasi real-time.

45. The method of Claim 40 wherein:

said database system comprises a buffer in cache memory; and

said normal operation comprises a buffer change operation.

46. The method of Claim 40 further comprising:

said database system during said normal operation, maintaining a count of physical writes from volatile memory to nonvolatile memory that actually occur;

wherein said maintaining is simulated during said simulating performance, to maintain a plurality of simulated counters of said physical writes that would have occurred for corresponding additional settings of MTTR being simulated.

47. The method of Claim 40 further comprising:

said database system during said normal operation, using a normal checkpointing queue; and

said simulating using a plurality of simulated checkpointing queues for corresponding additional settings of MTTR being simulated.

48. A computer-readable storage medium having stored thereon computer instructions that, when executed by a computer, cause the computer to, simulate effect on runtime performance of a database system, of a plurality of settings to control mean time to recover (MTTR) from a database failure, the computer instructions comprising instructions to:

use a current setting to limit MTTR of said database system for normal operation; and

simulate performance of said database system for a plurality of additional settings of MTTR, using data from said normal operation.

49. The computer-readable storage medium of Claim 48, wherein said data comprises actual database operating data from said database system.

50. The computer-readable storage medium of Claim 48, wherein:

said instructions to simulate are executed concurrently with said instructions to perform normal operation.

51. The computer-readable storage medium of Claim 50 wherein:

results of simulating are available in real-time.

52. The method of Claim 50 wherein:

results of simulating are available in quasi real-time.

53. The computer-readable storage medium of Claim 48 wherein:

said database system comprises a buffer in cache memory; and

said normal operation comprises a buffer change operation.

54. The computer-readable storage medium of Claim 48 further comprising instructions for:

said database system during said normal operation, to maintain a count of physical writes from volatile memory to nonvolatile memory that actually occur;

wherein said maintaining of the count is simulated by said instructions to simulate performance, to maintain a plurality of simulated counters of said physical writes that would have occurred for corresponding additional settings of MTTR being simulated.

55. The computer-readable storage medium of Claim 48 further comprising instructions for:

said database system during said normal operation, to use a normal checkpointing queue; and

wherein said instructions to simulate performance comprise instructions to use a plurality of simulated checkpointing queues for corresponding additional settings of MTTR being simulated.

56. A system, to simulate effect on runtime performance of a database system, of a plurality of settings to control mean time to recover (MTTR) from a database failure, the system comprising:

means for using a current setting to limit MTTR of said database system for normal operation; and

means for simulating performance of said database system for a plurality of additional settings of MTTR, using data from said normal operation.

57. The system of Claim 56, wherein said data comprises actual database operating data from said database system.

58. The system of Claim 56 wherein:

said means for simulating operates concurrently with said normal operation of said database system.

59. The system of Claim 58 wherein:

results of said simulating are available in real-time.

60. The method of Claim 58 wherein:

results of said simulating are available in quasi real-time.

61. The system of Claim 56 wherein:

said database system comprises a buffer in cache memory; and

said normal operation comprises a buffer change operation.

62. The system of Claim 56 further comprising:

means for maintaining a count of physical writes from volatile memory to nonvolatile memory that actually occur during said normal operation;

wherein said maintaining of the count is simulated in said means for simulating performance, to maintain a plurality of simulated counters of said physical writes that would have occurred for corresponding additional settings of MTTR being simulated.

63. The system of Claim 56 further comprising:

means for using a normal checkpointing queue during said normal operation;

wherein said means for simulating performance comprises means for using a plurality of simulated checkpointing queues for corresponding additional settings of MTTR being simulated.

A clean copy of the allowed claims is attached.

Reasons for Allowance

5. Claims 11, 13, 14, 16-22, 27-30, 34, 35 and 37-63 of the application are allowed over prior art of record.

6. The following is an Examiner's statement of reasons for the indication of allowable subject matter:

The closest prior art of record shows:

(1) checkpointing schemes to maintain copy of the primary database in main memory on a disk as a backup database; performing checkpointing incrementally and in parallel with normal transaction execution; support for fast restart after system crash by checkpointing continuously; storing in stable memory both UNDO and REDO logs produced by active transactions; reducing the amount of disk I/O by using recovery memory as a checkpointing buffer; evaluating the performance of the checkpointing schemes under a variety of workloads using database simulation using a queuing model to implement normal transaction processing and checkpointing procedure (**Lee et al.,** "checkpointing schemes for fast restart in main memory database systems", IEEE 1997);

(2) checkpointing and rollback recovery to keep the integrity of information and enhance reliability in database systems; determining optimum checkpointing policy that optimizes certain performance measure; checkpointing strategies include Poisson

checkpointing, fixed time interval between checkpoints, and fixed number of completed transactions between checkpoints; recovery period depends on checkpointing strategy selected; objective is to minimize the time spent in checkpointing and recovery; the optimum checkpointing strategy minimizes the mean response time of a transaction; using simulation to evaluate alternate checkpointing strategy (**Nocola et al.,** "Comparative analysis of different models of checkpointing and recovery", IEEE 1990); and

(3) checkpointing to ensure quick recovery in case of system crash; use of large database cache buffers in large database systems; checkpointing schemes scalable with the size of the buffer cache; checkpointing schemes that advance frequently to limit recovery time; scalability is achieved by organizing all modified buffers in ordered queues; frequent advancement of the database checkpoint made possible by incremental checkpointing; incremental checkpoints are continuous, low overhead checkpoints that write buffers as a background activity; incremental checkpointing continuously advances the database checkpoint i.e., the starting position in the redo log for crash recovery, resulting in dramatic reduction in recovery time; buffers are written from the head of the checkpoint queue in response to checkpoint request; the rate of buffer writes in incremental checkpointing is controlled to balance checkpoint writing time with recovery time requirements (**Joshi et al.,** "Checkpointing in Oracle", 24th VLDB conference, New York, 1998).

Additional state of the art reviewed and considered by the Examiner is found in U.S. Patent Application 2003/0084372; U.S. Patent 6,735,593; U.S. Patent 6,631,478; U.S. Patent Application 2004/0193658; U.S. Patent Application 2005/0131853.

None of these references taken either alone or in combination with the prior art of record discloses a computer implemented method that comprises a database system, to simulate effect of a recovery time required for the database system to recover from a database failure on runtime performance of the database system, specifically including:

"providing in the computer, in addition to a normal checkpoint queue used in the database system for normal operation, at least one simulated checkpoint queue;

wherein the simulated checkpoint queue is associated with a setting for recovery time whose effect on runtime performance of the database system is being simulated in the computer;

in response to detecting a change to a buffer in the normal checkpoint queue due to actual database transactions occurring within the database system under normal operating conditions, checking if the buffer is represented in the simulated checkpoint queue, and if the buffer is not represented in the simulated checkpoint queue, linking an element that represents the buffer to the tail of the simulated checkpoint queue;

providing a simulated write counter, the simulated write counter being associated with the setting for recovery time;

wherein the simulated write counter provides a count of a number of times any element is removed from the simulated checkpoint queue in response to a simulated write out of the respective buffer from volatile memory to nonvolatile memory;

determining if linking the element to the simulated checkpoint queue causes the simulated checkpoint queue to exceed a predetermined length; and

in response to determining that the simulated checkpoint queue exceeds the predetermined length, removing an element from the head of the simulated checkpoint queue and incrementing the simulated write counter".

None of these references taken either alone or in combination with the prior art of record discloses a computer-readable storage medium having stored thereon computer instructions that, when executed by a computer, cause the computer to simulate effect of a recovery time required for the database system to recover from a database failure on runtime performance of the database system, specifically including:

"provide in the computer, in addition to a normal checkpoint queue used in the database system for normal operation, at least one simulated checkpoint queue;

wherein the simulated checkpoint queue is associated with a setting for simulated mean time to recover (MTTR) whose effect on runtime performance of the database system is being simulated in the computer;

in response to detecting a change to a buffer in the normal checkpoint queue due to actual database transactions occurring within the database system under normal operating conditions, check if the buffer is represented in the simulated checkpoint queue, and if the buffer is not represented in the simulated checkpoint queue, link an element that represents the buffer to the tail of the simulated checkpoint queue;

provide a simulated write counter, the simulated write counter being associated with the simulated MTTR setting;

wherein the simulated write counter provides a count of a number of times any element is removed from the simulated checkpoint queue in response to a simulated write out of the respective buffer from volatile memory to nonvolatile memory;

determine if linking the element to the tail of the simulated checkpoint queue causes the simulated checkpoint queue to exceed a predetermined length; and

in response to determining that the simulated checkpoint queue exceeds the predetermined length, remove an element from the head of the simulated checkpoint queue and increment the simulated write counter".

None of these references taken either alone or in combination with the prior art of record discloses a system for simulating effect of a recovery time required for the database system to recover from a database failure on runtime performance of the database system, specifically including:

"a normal checkpoint queue used in the database system for normal operation, wherein the normal checkpoint queue comprises a plurality of buffers;

at least one simulated checkpoint queue;

wherein the simulated checkpoint queue is associated with a setting for simulated mean time to recover (MTTR) whose effect on runtime performance of the database system is being simulated in the computer; and

a simulated write counter maintained in the memory, the simulated write counter being associated with the simulated MTTR setting;

wherein the simulated write counter provides a count of the number of times an element is removed from the simulated checkpoint queue, in response to a simulated write out of a buffer from volatile memory and storing in nonvolatile memory;

wherein in response to detecting a change to a buffer in the normal checkpoint queue due to actual database transactions occurring within the database system under normal operating conditions, the instructions in memory check if the buffer is represented in the simulated checkpoint queue, and if the buffer is not represented in the simulated checkpoint queue, the instructions link an element that represents the buffer to the tail of the simulated checkpoint queue;

wherein the instructions in memory determine if linking the element to the simulated checkpoint queue causes the simulated checkpoint queue to exceed a predetermined length; and

in response to determining that the simulated checkpoint queue exceeds the predetermined length, the instructions in memory remove an element from the head of the simulated checkpoint queue and increment the simulated write counter".

None of these references taken either alone or in combination with the prior art of record discloses a computer-implemented method, computer-readable storage medium having computer instructions and a system to simulate effect on runtime performance of a database system, of a plurality of settings to control mean time to recover (MTTR) from a database failure, specifically including:

using a current setting to limit MTTR of said database system for normal operation; and

simulating performance of said database system for a plurality of additional settings of MTTR, using data from said normal operation.

7. Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

8. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Dr. Kandasamy Thangavelu whose telephone number is 571-272-3717. The examiner can normally be reached on Monday through Friday from 8:00 AM to 5:30 PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Leo Picard, can be reached on 571-272-3749. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to TC 2100 Group receptionist: 571-272-2100.

Paul L. Rodriguez 9/16/05

Primary Examiner
Art Unit 2125

K. Thangavelu
Art Unit 2123
September 15, 2005